Diffusion Models as Data Mining Tools

Ioannis Siglidis[1] Aleksander Holynski[2] Alexei A. Efros[2] Mathieu Aubry[1] Shiry Ginosar[2]
[1]Imagine, LIGM, Ecole Des Ponts [2]BAIR, University of California, Berkeley

Visual Mining

This looks **so** US! **Why?**

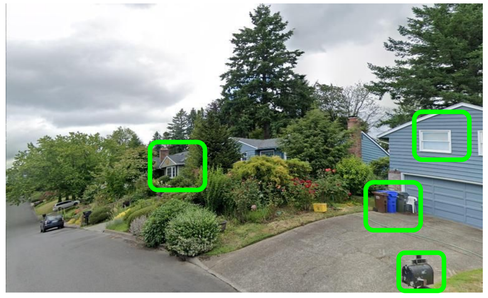

Traditional Mining

e.g. "What Makes Paris Look Like Paris", Doersch et al. 2013

Ours

Diffusion Model

Approach

Improved performance in conditional generation suggests typical elements.

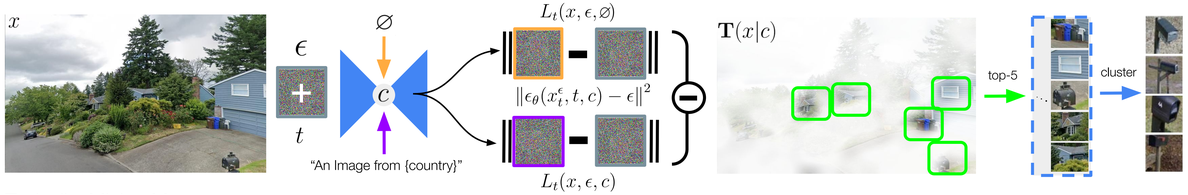

Typicality Mining Measure

$$\mathbf{T}(x|c) = \mathbb{E}_{\epsilon,t}[L_t(x, \epsilon, \varnothing) - L_t(x, \epsilon, c)]$$

"An element is ***typical*** if on average, the model can denoise it better using the conditioning label (c) than without information (∅)."

Baseline Comparison

our implementation*

Discovering Bias in Generation

Dataset was generated with c.f. guidance set to 7.5

Results: Visual Summaries

Our approach can scale to a large variety of datasets.

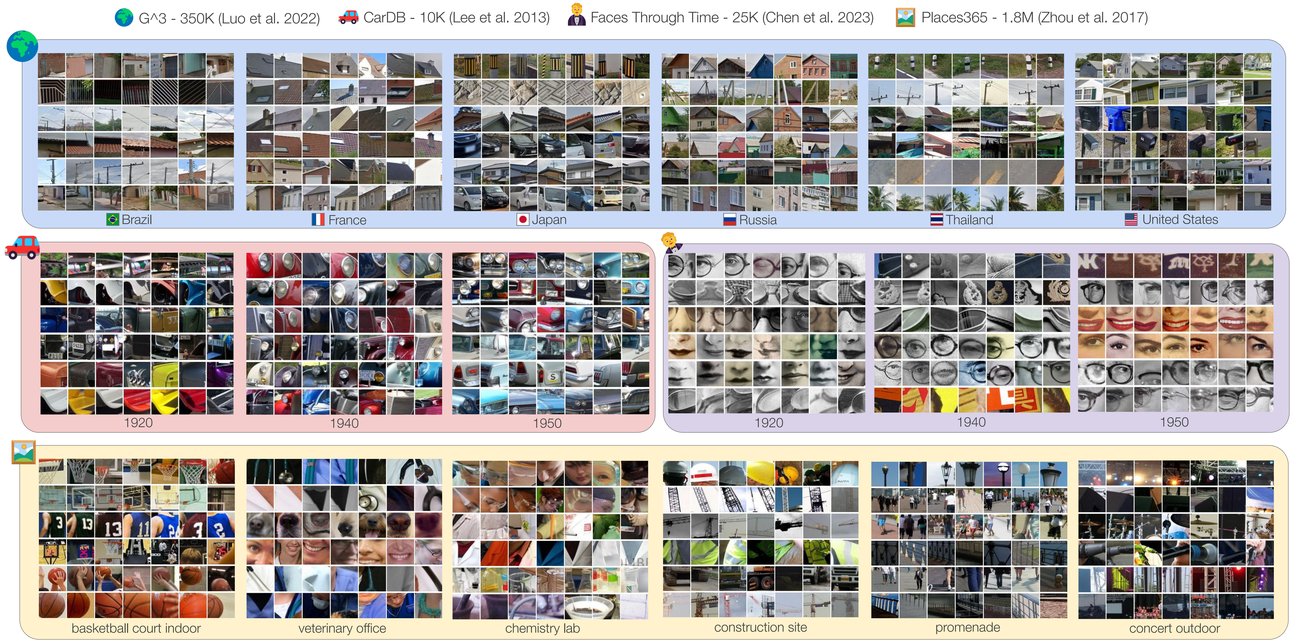

Application: Discovering co-typical Elements

Mining elements that are consistently typical across locations.

Some elements are typical to all places!

Scenes are rarely registered → Generate

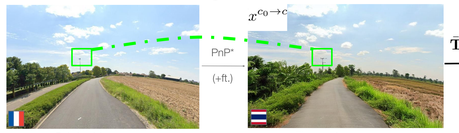

Extracting parallel patches according to $\bar{\mathbf{T}}(x) = \underset{c \in C}{\text{med}}[\mathbf{T}(x^{c_0 \to c}, c)]$ and then joint-clustering translated elements.